Articles   School of Accounting and Finance

2012

Hedging Effectiveness under Conditions of Asymmetry

Jim Hanly
Technological University Dublin, james.hanly@tudublin.ie

John Cotter
University College Dublin

Recommended Citation

Hedging Effectiveness under Conditions of Asymmetry

John Cotter and Jim Hanly

We examine whether hedging effectiveness is affected by asymmetry in the return distribution by applying tail specific metrics, for example, Value at Risk, to compare the hedging effectiveness of short and long hedgers. Comparisons are applied to a number of hedging strategies including OLS, and both symmetric and asymmetric GARCH models. We apply our analysis to a dataset consisting of S&P500 index cash and futures containing symmetric and asymmetric return distributions chosen *ex-post*. Our findings show that asymmetry reduces out-of-sample hedging performance and that significant differences occur in hedging performance between short and long hedgers.

1.   Introduction

A large literature has developed in the evaluation of futures based hedging strategies. The dominant hedging framework uses the variance as the risk measure and the Optimal Hedge Ratio (OHR) has, therefore, become synonymous with the minimum variance hedge ratio (see Lien and Tse, 2002, for a comprehensive survey). A shortcoming of the variance risk measure is that it cannot distinguish between positive and negative returns and therefore it does not allow for distribution asymmetries. Moreover it measures average risk only, and does not distinguish between specific parts of a return distribution such as the tail of the distribution.

In the literature on optimal hedging, two broad strands have emerged that address these issues. The first approach is the use of hedging estimation methods that seek to minimise some measure of risk other than the variance. These include Lower Partial Moments (LPM), Value at Risk (VaR) and Conditional Value at Risk (CVaR) and these are all tail-specific measures that can be applied separately to either side of a distribution. The second approach is the use of hedging estimation methods that allow for asymmetry in the return distribution. Of these, asymmetric extensions to the Generalised Autoregressive Conditional Heteroskedastic (GARCH) class of models have been applied. However no study has addressed the issue of hedging effectiveness under conditions of asymmetry, by combining these models, together with hedging effectiveness methods that measure tail probabilities, and we address this here.

First, we identify asymmetry in the bivariate distribution of cash and futures markets. In comparing hedging performance, it is common practice to examine and assess alternative models that estimate OHR's. We choose three of the most commonly applied. OLS is chosen as it has been the benchmark model to examine optimal futures hedging (see Ederington, 1979). We also use a Naive model that assumes the hedge ratio equals 1. To account for time variation

we use two different GARCH models, both symmetric and asymmetric, to examine the impact of asymmetry in both the volatility and covariance dynamics. Both of the GARCH models we use have been broadly applied in the literature (Brooks Henry and Persand, 2002) as have the OLS and Naïve models. This gives our study a solid platform on which to base our comparisons for symmetric and asymmetric data.

We apply tail specific hedging effectiveness measures to cash and futures containing symmetric and asymmetric return distributions chosen *ex-post*. This allows us to compare the hedging effectiveness of short and long hedgers under conditions of asymmetry. Our analysis could also be illustrated via simulation of alternative asymmetric distributions and associated parameters.[1] However, this would miss out on the real world application to financial data where it is standard that both excess skewness and excess kurtosis are evident in returns. Unlike excess kurtosis that indicates the presence of fat tails, there is no consensus on whether skewness of returns can be characterised as being either positive or negative, and nor do we have agreement on the magnitude of asymmetry. Also we are unable to identify a correct distribution for financial returns (with the exception that normality is usually rejected). To avoid having to identify a correct distribution with the correct parameter values, and especially the direction and magnitude of skewness, we use the far simpler approach of analysing real financial data and identifying timeframes where asymmetry is present and absent in order to examine hedging performance.

Combining these, we have four estimation methods and four performance metrics. This approach allows us to comprehensively examine hedging effectiveness of both short and long hedgers for both symmetric and asymmetric distributions, and to see whether there is a dominant OHR estimation method that emerges across a broad range of hedging effectiveness metrics.

Our most important finding is that the presence of skewness in the return distribution significantly reduces out-of-sample hedging effectiveness. Therefore hedges may underperform during stressful market conditions when they are most required. We also find significant differences between the hedging effectiveness of short and long hedgers for skewed returns. This implies that hedgers who fail to use tail specific hedging performance metrics may chose inefficient hedging strategies that result in being mishedged vis-à-vis their hedging objectives. Our results for the best hedging model are mixed but tend to favour using the OLS model. Based on these findings, it would appear that there is little to be gained in terms of hedging efficiency from the use of more complicated hedging strategies such as asymmetric GARCH models for both non-skewed and skewed distributions.

The remainder of the paper proceeds as follows. Section 2 outlines the methods used for estimating OHR's. Section 3 describes the metrics for measuring hedging effectiveness. Section 4 describes the data. Section 5 presents our empirical findings and Section 6 summarises and concludes the paper.

2. Methodology

The OHR is defined in the literature as the ratio that minimises the risk of the payoff of the hedged portfolio. The payoff of a hedged portfolio is given as:

$$+ r_s - \beta \ r_f \ \text{(short hedger)} \tag{1a}$$

$$- r_s + \beta \ r_f \ \text{(long hedger)} \tag{1b}$$

where r_s and r_f are returns on the cash and futures respectively, and β is the estimated OHR. We define a short (long) hedger as being long (short) the cash asset and short (long) the futures

asset. In this study we utilise four different methods for estimating OHR's benchmarked against a no-hedge strategy. The simplest model is a 1:1 or Naïve hedge ratio (h = 1) where each unit of the cash contract is hedged with equivalent units of the futures contract. The next method applied is an OLS HR, which is the slope coefficient of a regression of cash on futures returns. An OHR estimated by OLS associated with Ederington (1979) has been applied extensively in the literature. Cecchetti et al., (1988) argue that the OLS method is not optimal because it assumes that the OHR is constant, whereas time-varying volatility is the rule for financial time series. As the OHR depends on the conditional distribution of cash and futures returns, so too should the hedge ratio. We therefore use a rolling window OLS model to account for time varying effects. This is given as:

$$r_{st} = \alpha + \beta_t r_{ft} + \varepsilon_t \tag{2}$$

where r_{st} and r_{ft} are the cash and futures returns respectively for period t, ε_t is the disturbance term and β_t is the OHR. This can also be expressed as:

$$\beta_t = \frac{h_{sft}}{h_{ft}} \tag{3}$$

where h_{ft} denotes the variance of futures returns and h_{sft} is the covariance between cash and futures returns. We also use two additional GARCH based estimation methods that allow the OHR to be time varying.

2.1   *The Symmetric Diagonal VECH GARCH Model (SDVECH)*

We use the Diagonal Vech (1,1) GARCH model proposed by Bollerslev, Engle and Wooldridge (1988). This model imposes a symmetric response on the variance and is useful for comparison of hedging estimation and performance as it has been extensively applied in the hedging literature. This model is specified as follows:

$$r_{st} = \mu_{st} + \varepsilon_{st} \quad r_{ft} = \mu_{ft} + \varepsilon_{ft}, \qquad \begin{bmatrix} \varepsilon_{st} \\ \varepsilon_{ft} \end{bmatrix} \middle| \Omega_{t-1} \sim N\left(0, h_t\right) \tag{4}$$

$$h_{11,t+1} = w_s + a_{ss}\varepsilon_{s,t}\varepsilon_{s,t} + b_{ff,t}h_{ff,t} \tag{5}$$

$$h_{22,t+1} = w_f + a_{ff}\varepsilon_{f,t}\varepsilon_{f,t} + b_{ff,t}h_{ff,t} \tag{6}$$

$$h_{12,t+1} = w_{sf} + a_{sf}\varepsilon_{s,t}\varepsilon_{f,t} + b_{sf,t}h_{sf,t} \tag{7}$$

where r_{st}, r_{ft} are the returns on cash and futures respectively, $\varepsilon_{s,t}, \varepsilon_{f,t}$ are the residuals,

$h_{ss,t+1}, h_{ff,t+1}$ denotes the variance of cash and futures and $h_{sf,t+1}$ is the covariance, w is a 3x1

parameter vector, and a and b are 3 x 3 parameter matrices. The matrices a and b are restricted to

be diagonal implying that the conditional variance of the cash returns depends only on past

values of itself and past values of the squared innovations in the cash returns. The conditional

variance of the futures returns and the conditional covariance between cash and futures returns

have similar structures. Because of the diagonal restriction we use only the upper triangular

portion of the variance covariance matrix, the model is therefore parsimonious, with only nine

parameters in the conditional variance-covariance structure of the Diagonal VECH model to be

estimated. This is subject to the requirement that the variance-covariance matrix is positive

definite for all values of ε_t in order to generate positive hedge ratios.

2.2   *The Asymmetric Diagonal VECH GARCH Model (ASDVECH)*

The second GARCH model that we use is an asymmetric extension of the SDVECH model. The

asymmetric model is able to capture asymmetries both within and between cash and futures

markets as it allows volatility to respond differently to negative and positive returns. The

asymmetric GARCH model builds on the univariate asymmetric GARCH model of Glosten,

Jagannathan and Runkle (GJR) (1993) and allows the variance to respond differently to positive

and negative return innovations through the use of an additional term designed to capture

5

asymmetries. The multivariate model is similar to that used in de Goeij and Marquering (2004). It differs from the SDVECH model by changing the equations for the variances of the cash and futures given in (5) and (6) as follows:

$$h_{11,t+1} = w_s + a_{ss}\varepsilon_{s,t}\varepsilon_{s,t} + b_{ss,t}h_{ss,t} + a^1{}_{ss}\varepsilon_{s,t},\varepsilon_{s,t}I_{t-1} \tag{8}$$

$$h_{11,t+1} = w_f + a_{ff}\varepsilon_{f,t}\varepsilon_{f,t} + b_{ff,t}h_{ff,t} + a^1{}_{ff}\varepsilon_{f,t},\varepsilon_{f,t}I_{t-1} \tag{9}$$

Also the covariances are modelled to allow for asymmetric shocks using:

$$h_{sf,t+1} = w_{sf} + b_{sf}h_{sf,t} + a^1{}_{sf,t}\varepsilon_{s,t}\varepsilon_{f,t} + a^2{}_{sf}I_{\varepsilon s,t},\varepsilon_{s,t}I_{\varepsilon f,t}\varepsilon_{f,t} + a^3{}_{sf}I_{\varepsilon s,t}\varepsilon_{s,t}(1-I_{\varepsilon f,t})\varepsilon_{f,t}$$
$$+ a^4{}_{sf}(1-I_{s,t})\varepsilon_{s,t}I_{\varepsilon f,t}\varepsilon_{f,t} \tag{10}$$

Where $I_{\varepsilon k,t}$ is an indicator variable equal to one if $\varepsilon_{k,t}, < 0$, and zero otherwise where $k = 1,2$ such that the space can be partitioned into four quadrants in the $\varepsilon_1, \varepsilon_2$ plain to allow for positive and negative shocks.

GARCH models have been broadly applied in the hedging literature to estimate time varying optimal hedge ratios (see Kavussanos and Visvikis, 2008). However, the performance of these models has been mixed. Over short time horizons and in-sample they have performed well (see Sultan and Hasan, 2008); however, over longer hedging horizons and out-of-sample, their performance has been poor. Brooks et al, (2002) use an asymmetric model to compare the hedging performance of a symmetric and asymmetric GARCH model. They find that accounting for asymmetry may yield marginal improvements in hedging efficiency in-sample but not for out-of-sample hedging. In contrast, Lien and Yang (2006) examine the effects of asymmetry in the spread between currency spot and futures on hedging performance. They find that a GARCH model incorporating asymmetric effects outperforms a symmetric model for both in–sample and out-of-sample comparisons. However, they confine their examination of hedging effectiveness to the variance metric in contrast to our use of a number of measures, and importantly, tail-specific

ones. Switzer and El-Khoury (2007) also examine hedging using the GJR GARCH model applied to weekly oil data. However their approach differs from ours in that they define asymmetry on a univariate dimension whereas the impact of asymmetry in hedging should be examined in the bivariate dimension. We now turn to hedging effectiveness in the next section.

3.   Hedging Effectiveness

In this paper we use four different performance metrics that have been applied in the hedging and risk management literatures. Three of these, LPM, VaR and CVaR were chosen as they are tail specific measures and can examine risk separately for both sides of the distribution of hedged portfolios. They are, therefore, particularly suitable for a study into the effects of asymmetry on hedging effectiveness. They can also measure performance for short and long hedgers separately. As a benchmark comparison we also included the two-sided risk measure, the variance in our performance evaluation. For each risk measure, performance is measured by the percentage reduction compared to an unhedged portfolio.[2]

The LPM is a hedging effectiveness metric that we use to allow for asymmetric data as it can distinguish between the left and right tails of a distribution. The LPM measures the probability of falling below a pre-specified target return (Bawa, 1975). The LPM of order n around τ is defined as

$$\text{LPM}_n(\tau; R) \quad = E\left\{(\max[0, \tau - R])^n\right\} \equiv \int_{-\infty}^{\tau} (\tau - R)^n \, dF(R) \quad (11)$$

where F(R) is the cumulative distribution function of the investment return R[3], and τ is the target return parameter. The value of τ will depend on the level of return or loss that is acceptable to the investor. Some values of τ that may be considered are, zero or the risk free rate of interest. The parameter n is the weighting applied to shortfalls from the target return. The more risk averse an

investor the higher the weight (*n*) that would be attached (Fishburn, 1977). We can form a complete set of downside risk measures by changing the τ and *n* parameters to reflect the position and risk preferences of different types of hedger. The more risk averse investors may set τ as the disaster level of return and have utilities that would reflect an LPM with *n*=2, 3… We choose $\tau = 0$ and $n = 3$ describing a risk-averse investor.[4]

Because of its focus on tail specific risk, he LPM serves as an intuitive risk measure that is in line with the risk preferences of many investors (Lee and Rao, 1988). Price and Nantell (1982) in a comparison of LPM and variance based measures of hedging effectiveness, find that they provide equivalent measures for normal distributions but that the LPM differentiates for distributions that are asymmetric, as it is a tail specific measure that measures the left and right tail probabilities independently. Studies have compared the hedging effectiveness of short and long hedgers using the LPM methodology (Demirer and Lien, 2003). The results indicate that hedging effectiveness for long hedgers differs from that of short hedgers; with long hedgers deriving more benefit from hedging in terms of risk reduction as measured by reduced LPM's.

The second hedging effectiveness metric is VaR. This is the loss level of a portfolio over a certain period that will not be exceeded with a specified probability. VaR has two parameters, the time horizon (N) and the confidence level (x). Generally VaR is the (100-x)[th] percentile of portfolio returns over the next N days. We calculate VaR using $x = 99$ and $N = 1$.

A shortcoming in VaR is that it is not a coherent measure of risk, as it is not subadditive[5]. Also in practice two portfolios may have the same VaR but different potential losses as it does not account for losses beyond the (100-x)[th] percentile. We address this by estimating an additional performance metric; CVaR, a coherent measure of risk. CVaR is the expected loss conditional that we have exceeded the VaR. It is given as:

$$CVaR = E[L|L > VaR] \tag{12}$$

This measures the expected value of our losses, L, in excess of the VaR. CVaR is attractive to hedgers because it estimates not only of the probability of a loss, but also the magnitude of a possible loss. In calculating CVaR we use the 1% confidence level which gives us the expected loss beyond the 1% VaR.

Few studies in the hedging literature have applied either VaR or CVaR as measures of hedging effectiveness in combination with GARCH models. However, Giot and Laurent (2003) use both the VAR and CVaR measures to examine the risk of short and long trading positions over a one day time horizon, estimating volatility with both symmetric and asymmetric GARCH approaches. They show that symmetric models underperform models that account for asymmetry; however, their analysis is only applied to unhedged positions.

4.   Data

Our data consist of daily cash and futures closing prices of the S&P500 Equity Index. To examine equity hedging, we chose the S&P500 index because of its economic importance and because it has been extensively used in the hedging literature as a benchmark in examining hedging effectiveness. All data was obtained from Datastream and daily returns were calculated as the first differenced logarithmic prices. Continuous series were formed by using the nearby contract with rollover driven by choosing the highest volume contract.

As we are examining the influence of asymmetry on hedging effectiveness, we require data that exhibits both skewed and non-skewed characteristics in order to facilitate a comparison. These considerations motivated our choice of sample periods for the S&P500 data. Our initial sample is daily logarithmic returns from January 1, 2000 through December 31, 2008. Because hedging is concerned with the bivariate return distribution of cash and futures, the measure of asymmetry we use is the skewness of the hedged portfolio comprising one cash and one futures

9

contract. The second criterion we use is that both in-sample and out-of-sample periods must be consistently symmetrically or non-symmetrically distributed to ensure consistency for the in-sample and out-of-sample performance results. Following these criteria, we picked a sample of daily cash and futures prices for the period January 2004 – December 2006. From these sample periods we were able to extract two separate equal sized datasets, one skewed and one non-skewed. The symmetric period runs from January 2004 – December 2004, while the asymmetric period runs from January 2006 – December 2006.

For each period, the first 160 observations are used for the in-sample estimation of the hedging models. The remaining 100 observations were used to facilitate out-of-sample comparisons for model evaluation. The characteristics of the data are now examined.

INSERT TABLE 1 HERE

Summary statistics for the hedged portfolio are displayed in Table 1. We see that the characteristics of the return distributions of the two series are markedly different. Returns in 2004 are symmetric - insignificant skewness (-0.01). In contrast, returns in 2006 are asymmetric - significant negative skewness (-1.51). Furthermore, the 2004 series is normal as measured by an insignificant Bera Jargue statistic while the 2006 series is non-normal. The data were checked for stationarity using Dickey Fuller unit root tests. As expected, the series is stationary. Stationarity is important as a non-stationary series may lead to spurious regressions and invalidate the estimated optimal hedge ratios.

5.   **Empirical Results**

The in-sample hedge ratios are graphically represented in Figure 1 together with the associated summary statistics in Table 2. The results of fitting the hedging estimation models are quite standard and are not reproduced here (available on request).

INSERT FIGURE 1 HERE

INSERT TABLE 2 HERE

A quick glance at Figure 1 shows that there are differences in the characteristics of the OHR's for the symmetric and asymmetric series. We see that the SDVECH and ASDVECH models are very similar for the symmetric series. This is confirmed if we examine the summary statistics of the two hedging models in Table 2. If we compare the mean hedge ratios of the two models using standard t-tests we find no significant differences between the SDVECH and ASDVECH models (t-stat 0.08). This result isn't surprising given the symmetric nature of the underlying distribution of the series for 2004. When we examine the asymmetric series, however, we find significant differences between the mean OHR for the SDVECH model and the ASDVECH model (t-stat 3.12), and supported in Figure 1. Therefore the choice of hedging model is more relevant where there is asymmetry in the bivariate distribution of cash and futures.

Turning next to hedging performance, Table 3 presents in-sample and out-of-sample results for hedging effectiveness for both symmetric and asymmetric data. Examining first the in-sample results, these show that hedging is effective at reducing risk as measured by each of the hedging effectiveness metrics. Risk reductions from hedging range from about 60% to almost 99% depending on the risk metric and the model used. For example, if we look at the symmetric series, we find reductions of the order of 94% - 98% in the Variance and LPM respectively, whereas reductions in both the VaR and CVaR are in the region of 73% - 77% for the best performing model.

INSERT TABLE 3 HERE

These results may be driven by VaR and CVaR metrics modelling extreme tail events whereas the LPM is a more general metric that doesn't pick up the most extreme outliers in the same way.[6] It is clear that hedging may be more limited in reducing extreme losses as measured

by the VaR or CVaR. A key point is that hedges may not be as effective at reducing exposure to tail risk as compared with more general measures of volatility such as the variance. In hedging terms, this means that investors may face the risk that their hedges will not be as effective during stressful markets conditions when they are most needed.

We also make a statistical comparison of the hedging effectiveness of short and long hedgers using Efrons (1979) bootstrap methodology, by employing t-tests of the differences between short and long hedgers based on the point estimates of our results. This approach is also adopted in tests of model hedging effectiveness and allows us to make statistical as well as economic inferences from our results[7]. In-Sample, the differences in hedging effectiveness between short and long hedgers for asymmetric data are significant at the 1% level in 75% of cases, whereas for the symmetric sample they are significant in only 33% of cases. The differences between short and long hedgers even for the symmetric period demonstrates the importance of using tail specific hedging effectiveness metrics, irrespective of the characteristics of the return distribution.

We now turn to the out-of-sample results. When we examine hedging performance across each performance metric and for each hedging model, we find that the asymmetric period yields lower hedging performance as compared with the symmetric period, by an average of 22%. The performance is even worse if we look at tail measures of risk, the VaR and the CVaR. Taking the short hedgers as an example, the average reduction in VaR across all hedging models is 74% as compared with just 20% for the asymmetric period. The results for long hedgers are similar. Thus, we conclude that asymmetry in the joint distribution of cash and futures, reduces hedging effectiveness, irrespective of the risk measure used, or the hedging model employed.

Using the tail specific performance metrics, we also compare the out-of-sample hedging performance of short and long hedgers, for both symmetric and asymmetric distributions. For the asymmetric sample, we find evidence of significant differences between short and long hedgers in every single case. For example, long hedgers significantly outperform short hedgers, both statistically and economically, based on a comparison of VaR and CVaR for all hedging models. To illustrate the economic differences, a long hedger who uses an OLS model to hedge an exposure of $1,000,000 will reduce the VaR from $120,190 to $36,340, a reduction of 69.8%. A short hedger by comparison will reduce their exposure from $99,780 to $76,100, a reduction of 23.7%. We also find that for the symmetric sample there are significant differences between short and long hedgers in 58% of cases, although the differences are not as pronounced as for the asymmetric data.

We next examine hedging model performance. Table 4 presents absolute figures for the hedging effectiveness measures of each of the hedging models, for both in-sample and out-of-sample hedges. The lower values represent better hedging performance. For example, the variance of 0.5219 for the symmetric hedge represents the risk associated with an unhedged position, whereas the SDVECH model is the best hedging model as it yields the lowest risk with a variance of 0.0249. Looking first at the in-sample results, we see that the best hedging model depends on the hedging performance metric. For example, using symmetric data for long hedgers, the SDVECH model performs best in terms of the Variance (0.0249) VaR (0.364) and the CVaR (0.426), while the ASDVECH hedge is the best performing model using LPM (0.0037). However, while the ASDVECH model performs well in a number of cases, its failure to outperform the SDVECH model even when the data are asymmetric would seem to indicate that the extra complexity of this model does not result in better performance. This supports the findings of Brooks et al (2002) who find that the symmetric GARCH model tends to perform as

well as the asymmetric GARCH model in an out-of-sample setting. The out-of sample results show that the OLS and Naive models are the better performers for the symmetric data. For the asymmetric period, the OLS model is the best performer in 7 out of 8 cases.

While these findings highlight the best hedging models for a given scenario, a key issue is whether there are significant differences in the performance of the different hedging models. To test this we compared the performance of the different hedging models, again employing Efrons (1979) bootstrap methodology with the results presented in Table 4. We find significant differences in the hedging effectiveness of the different models in only 39% of cases in-sample, and 20% of cases out-of-sample. We also find that the absolute differences in performance between models are small and not economically significant. This demonstrates that an OLS hedge can provide as good or better hedging performance than more complicated models based on time-varying GARCH specifications, especially in an out-of-sample setting. This further supports the findings in the literature which show that OLS tends to perform as well or better than GARCH models. (Cotter and Hanly, 2006)

6.   Conclusion

This paper compares the hedging effectiveness of S&P500 equity index futures for both symmetric and asymmetric distributions. We also compare the hedging effectiveness of short and long hedgers using a variety of tail specific performance measures, including VaR and CVaR. We find that out-of-sample hedging effectiveness is significantly reduced by the presence of skewness in the bivariate return distribution of cash and futures. This is an important finding as it means that hedging may not be as effective during asymmetric return periods. We also find larger differences in hedging performance between the short and long hedgers for the

14

asymmetric distribution as compared with a symmetric distribution. Therefore the use of one-sided hedging performance measures that are consistent with modern risk management techniques such as VaR and CVaR is to be recommended. The traditional variance reduction criterion is not adequate, and will provide inaccurate measures of risk for different types of hedgers, both for symmetric and especially asymmetric distributions.

Results for the best overall hedging estimation model are mixed. The OLS model provides good hedging performance across all measures of hedging effectiveness, for both short and long hedgers. The GARCH models perform well in-sample, but the differences in performance compared with the OLS model are not significant. We also find that employing a GARCH model that allows for asymmetries yields no significant improvement in performance over the standard GARCH model.

References

Artzner, P., Delbaen, F., Eber, J., & Heath, D. (1999). Coherent measures of risk. Mathematical Finance, 9, 203 - 228.

Bawa, S. (1975). Optimal rules for ordering uncertain prospects. Journal of Financial Economics, 2, 95-121.

Bollerslev, T., Engle, R., & Wooldridge, J. (1988). A capital asset pricing model with time-varying covariances. Journal of Political Economy 96, 116 - 131.

Brooks, C., Henry, O., & Persand, G. (2002). The effects of asymmetries on optimal hedge ratios. Journal of Business, 75, 333 – 352.

Cecchetti, S., Cumby, R., & Figlewski, S. (1988). Estimation of the optimal futures hedge. Review of Economics and Statistics, 70, 623 – 630.

Cotter, J., & Hanly, J. (2006). Re-examining hedging performance. Journal of Futures Markets, 26, 657-676.

de Goeij, P., & Marquering, W. (2004). Modeling the conditional covariance between stock and bond returns: a multivariate GARCH approach. Journal of Financial Econometrics 4, 531–564.

Demirer, R., & Lien, D. (2003). Downside risk for short and long hedgers. International Review of Economics and Finance. 12, 25 – 44.

Ederington, L. (1979). The hedging performance of the new futures markets. Journal of Finance, 34, 157 – 170.

Efron, B. (1979). Bootstrap methods: Another look at the Jack-knife. The Annals of Statistics, 7, 1–26.

Fishburn, P. (1977). Mean-risk analysis with risk associated with below-target returns. The American Economic Review, 67, 116–126.

Giot, P., & Laurent, S. (2003). Value-at-Risk for long and short trading positions. Journal of Applied Econometrics, 18, 641-663.

Glosten, L., Jagannathan, R., & Runkle, D. (1993). On the relationship between the expected value and the volatility of the normal excess return on stocks. Journal of Finance, 48, 1779-1801.

Kavussanos, M. & Visvikis, I. (2008). Hedging effectiveness of the Athens stock index futures contracts. The European Journal of Finance, 14, 243–270.

Lee, W. & Rao, R. (1988). Mean lower partial moment valuation and lognormally distributed returns. Management Science 34, 446–453.

Lien, D., & Tse, Y. (2002). Some recent developments in futures hedging. Journal of Economic Surveys, 16, 357 – 396.

Lien, D., & Yang, L. (2006). Spot-futures spread, time-varying correlation, and hedging with currency futures. Journal of Futures Markets, 26, 357 – 396.

Price, K., Price, B., & Nantell, T. (1982). Variance and lower partial moment measures of systematic risk: Some analytical and empirical results. Journal of Finance, 37, 843- 855.

Sultan. A. and Hasan, B. (2008). The effectiveness of dynamic hedging: evidence from selected European stock index futures. The European Journal of Finance, 14, 469–488.

Switzer, L., & El-Khoury, M. (2007). Extreme volatility, speculative efficiency and the hedging effectiveness of the oil futures markets. Journal of Futures Markets, 27, 61 – 84.

Notes

1. Simulation, although not pursued here, offers a number of advantages including allowing us to address a range of complexities such as non-linearities in the relationship between spot and futures.

2. The percentage reduction in the relevant performance measure is generally compared with a no-hedge position. However there may be some cases where a no-hedge position does not yield the worst hedging performance, whereby the comparison is than made against the hedging model that yields the worst hedging performance.

3. Rather than assuming a particular distribution we use the empirical distribution of returns to estimate the LPM

4. We used a number of values for n, n=1 up to n=5 which yielded similar results and are not reported.

5. Subadditivity implies that the risk of two positions when added together is never greater than the sum of the risks of the two individual positions (see, for example, Artzner et al, 1999).

6. While each of these measures is one sided, the LPM with a target rate set t=0 will include all observations less than 0, whereas both the VaR and CVaR calculated at the 1% interval will include only extreme observations located in the left or right tails of the distribution.

7. The returns of the hedged portfolios as compiled using equations 1a and 1b were bootstrapped by resampling with replacement from the returns.100 simulations were used allowing for the construction of confidence intervals around each point estimate.

Table 1. Summary Statistics - Hedged Portfolio

	Mean x10^{-4}	Min	Max	Std Dev	Skewness	Kurtosis	B-J	Stationarity
Symmetric 2004	-0.080	-0.005	0.006	0.002	-0.01	0.75*	6.1	-28.8*
Asymmetric 2006	-0.090	-0.009	0.005	0.002	-1.51*	7.50*	707.9*	-19.6*

Notes: This table presents summary statistics for the Naive hedged portfolio for the S&P500 which comprises one cash and one futures contract. The skewness statistic measures asymmetry where zero would indicate a symmetric distribution. The kurtosis statistic measures the shape of a distribution where a value of zero would indicate a normal distribution. The Bera-Jarque B-J statistic combines skewness and kurtosis in comparison to normality. Stationarity is tested using the Dickey-Fuller unit root test. We characterise the 2004 series as symmetric given the insignificant skewness figure of -0.01. Similarly we characterise the 2006 series as asymmetric with a significant skewness figure of -1.51. *Denotes Significant at the 1% level.

Table 2. Summary Statistics of Optimal Hedge Ratios

	Symmetric 2004		Asymmetric 2006	
Hedging Model	SDVECH	ASDVECH	SDVECH	ASDVECH
Mean	0.987	0.986	0.950*	0.969
Std Dev	0.032	0.031	0.046	0.065
Min	0.831	0.845	0.774	0.711
Max	1.078	1.082	0.995	1.092
Stationarity	-8.60*	-8.28*	-6.94*	-4.58*

Notes: This table presents summary statistics for the estimated time-varying optimal hedge ratios for each hedging model. A statistical comparison is made between the mean of the SDVECH and ASDVECH hedge ratios for each period based on a difference in means t-test. Results show that for the asymmetric period there is a significant difference between the SDVECH OHR (0.95) and the ASDVECH OHR (0.969). Stationarity is tested using the Dickey-Fuller unit root test. * Denotes significant at the 1% level.

Table 3. Hedging Effectiveness – Short Vs Long Comparison

|	Variance	LPM	VaR	CVaR	Variance	LPM	VaR	CVaR
	Panel A: SHORT HEDGERS				**Panel B: LONG HEDGERS**			
				IN-SAMPLE				
Symmetric 2004								
No Hedge	0	0	0	0	0	0	0	0
Naïve	93.66	98.01	72.18	70.50*	93.66	98.47	72.61	66.58
OLS	93.77	98.19	72.02	71.69*	93.77	98.53	72.88	69.44
SDVECH	95.22	98.91	74.42	72.99	95.22	98.87	77.06*	73.54
ASDVECH	95.12	98.84	74.41	72.93	95.12	98.88	77.03*	73.48
Asymmetric 2006								
No Hedge	0	0	0	0	0	0	0	0
Naïve	94.32	99.05	65.62	61.33	94.32	97.42	82.88*	79.97*
OLS	94.53	99.03	68.43	65.22	94.53	97.87	83.77*	79.17*
SDVECH	94.05	98.45	70.76	69.06	94.05	98.18	74.38*	71.80*
ASDVECH	94.14	99.15*	65.62	60.07	94.14	96.76	84.29*	78.32*
				OUT-OF-SAMPLE				
Symmetric 2004								
No Hedge	0	0	0	0	0	0	0	0
Naïve	95.57	99.48	73.28	78.06	95.57	98.55	82.28*	81.53*
OLS	95.37	99.34	74.40	79.42	95.37	98.67	80.61*	79.23
SDVECH	95.47	99.43	74.05	77.42	95.47	98.57	81.55*	81.72*
ASDVECH	95.44	99.40	74.24	77.21	95.44	98.61	80.58*	80.10*
Asymmetric 2006								
No Hedge	0	0	0	0	0	0	0	0
Naïve	85.03	98.10*	14.77	31.26	85.03	71.15	69.25*	62.94*
OLS	85.43	97.92*	23.73	33.89	85.43	75.67	69.76*	63.16*
SDVECH	85.41	97.93*	22.15	33.67	85.41	75.14	69.62*	63.09*
ASDVECH	85.16	97.97*	19.73	32.07	85.16	73.37	69.51*	63.01*

Notes: This table presents the in-sample and out-of-sample hedging performance for the S&P500 returns for both the asymmetric and symmetric datasets. Results reported give the percentage reduction in the performance metric from the hedged model as compared with the worst hedged position. For example, short hedging the S&P500 contract for the asymmetric dataset with the OLS model yields a 94.53% in-sample reduction in the variance as compared with a No-Hedge strategy. Statistical comparisons are made between the performance of short and long hedgers on a metric by metric basis using Efron's 1979 bootstrap technique. * Indicates that the percentage reduction in the risk metric is significantly better comparing long to short hedges at the 1% confidence level.

Table 4. Statistical Comparison of Hedging Model Performance

	Variance x10^{-3}	LPM x10^{-5}	VaR x10^{-2}	CVaR x10^{-2}	Variance x10^{-3}	LPM x10^{-5}	VaR x10^{-2}	CVaR x10^{-2}
	Panel A: SHORT HEDGERS				Panel B: LONG HEDGERS			
	IN-SAMPLE							
Symmetric 2004								
None	0.5219*	0.2617*	1.5632*	1.607*	0.5219*	0.3347*	1.5892*	1.610*
Naïve	0.0331*	0.0052*	0.4349*	0.474*	0.0331*	0.0051*	0.4353*	0.538*
OLS	0.0325*	0.0047*	0.4374*	0.455*	0.0325*	0.0049*	0.4310*	0.492*
SDVECH	0.0249[a]	0.0029[a]	0.3998[a]	0.434[a]	0.0249[a]	0.0038	0.3645[a]	0.426[a]
ASDVECH	0.0255	0.0030	0.4001	0.435	0.0255	0.0037[a]	0.3651	0.427
Asymmetric 2006								
None	0.4904*	0.3533*	1.7384*	1.823*	0.4904*	0.2639*	1.9462*	2.117*
Naïve	0.0279	0.0034*	0.5977*	0.705*	0.0279	0.0068*	0.3332*	0.424[a]
OLS	0.0268[a]	0.0034*	0.5488*	0.634*	0.0268[a]	0.0056	0.3159	0.441*
SDVECH	0.0292*	0.0055*	0.5083[a]	0.564[a]	0.0292*	0.0048[a]	0.4986*	0.597*
ASDVECH	0.0288*	0.0030[a]	0.5977*	0.728*	0.0288*	0.0085*	0.3058[a]	0.459*
	OUT-OF-SAMPLE							
Symmetric 2004								
None	0.3868*	0.2758*	1.1252*	1.404*	0.3868*	0.1315*	1.5076*	1.603*
Naïve	0.0171[a]	0.0014[a]	0.3007	0.308*	0.0171[a]	0.0019	0.2671[a]	0.296
OLS	0.0179	0.0018	0.2880[a]	0.289[a]	0.0179	0.0018[a]	0.2923*	0.333*
SDVECH	0.0175	0.0016	0.2920	0.317*	0.0175	0.0019	0.2782	0.293[a]
ASDVECH	0.0177	0.0017	0.2898	0.320*	0.0177	0.0018	0.2928	0.319*
Asymmetric 2006								
None	0.2136*	0.1333*	0.9978*	1.369*	0.2136*	0.0578*	1.2019*	1.360*
Naïve	0.0320	0.0025[a]	0.8504*	0.941	0.0320	0.0167*	0.3696	0.504
OLS	0.0311[a]	0.0028*	0.7610[a]	0.905[a]	0.0311[a]	0.0140[a]	0.3634[a]	0.501[a]
SDVECH	0.0312	0.0028	0.7768	0.908	0.0312	0.0144	0.3651	0.502
ASDVECH	0.0317	0.0027	0.8009	0.930	0.0317	0.0154	0.3665	0.504

Notes: This Table presents the in-sample and out of sample hedged portfolio statistics upon which we base our performance measures. The best performing model is the model that yields the lowest value for each risk measure and is denoted by a. For example, the OLS model yields the lowest variance of 0.268 when hedging the in-sample asymmetric data. Statistical comparisons are made for each hedging model against the best performing model using Efron's 1979 technique. For example, taking the out-of-sample asymmetric dataset for a short hedger, we can see that there is a significant difference between the variance of both the No Hedge model 0.2136 and the best performing OLS model 0.0311, however there is no significant differences between the OLS as compared with the Naïve, SDVECH or ASDVECH models * Denotes a better performance being recorded for the best performing benchmark relative to that measure at the 1% significance level.

Figure 1. In-sample OHR's for the S&P500 series.
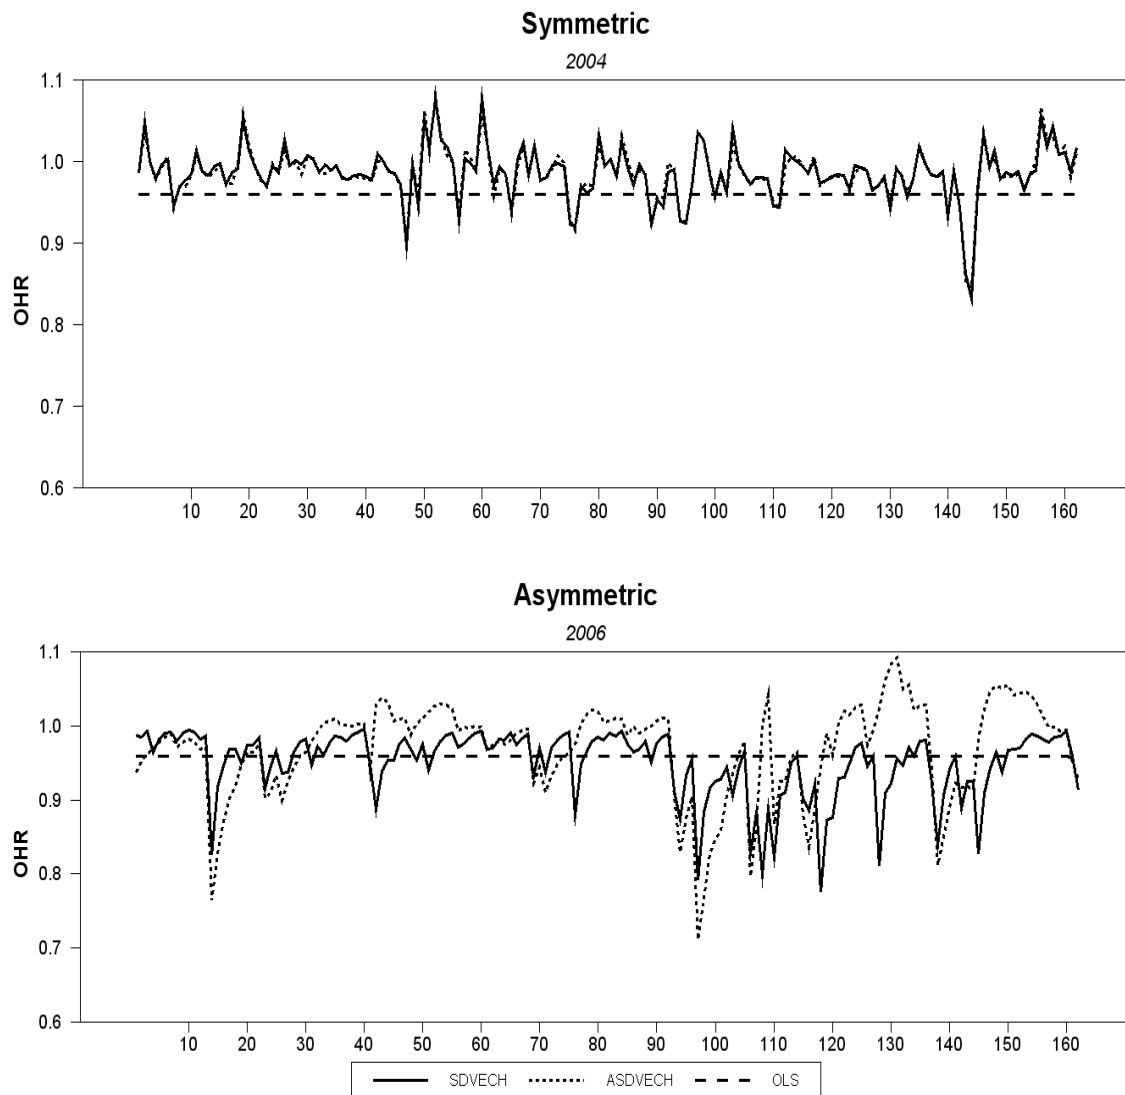
Notes: Three OHR's are presented for both the symmetric and asymmetric distributions which are 2004 and 2006 respectively. The OHR's are the OLS, symmetric GARCH SDVECH and asymmetric GARCH ASDVECH models. Both of the GARCH models yield similar hedges for the symmetric period whereas for the asymmetric period there is a significant difference between the SDVECH and ASDVECH models.